

Independence Results around Constructive ZF

Robert S. Lubarsky

October 5, 2015

1 Introduction

When making the transition from a classical to an intuitionistic system, one is compelled to reconsider the axioms chosen. For instance, it would be natural to re-formulate those that imply Excluded Middle. In the context of set theory, typically the Axiom of Foundation is substituted by the Axiom of Set Induction, the two being equivalent classically but the latter not yielding Excluded Middle intuitionistically. Even axioms that do not return classical logic must be re-evaluated, as classically equivalent variants may not remain so in a different logic. An example of this is Collection, which, over the other ZF axioms, is equivalent to Replacement and to Reflection, but not over intuitionistic ZF (see [4]).

A different reason to re-evaluate axioms is to take more seriously the program of constructivism. IZF is basically ZF with classical logic replaced by intuitionistic. Since, though, the axioms of ZF were never meant to be constructive, they're still not even with a different logic. This issue was addressed by Myhill (see [7]), who re-formulated by the axioms and framework of ZF to come up with CST (Constructive Set Theory). Among the changes he made was to eliminate the Power Set Axiom as non-constructive; as an alternative, he included Exponentiation, the hypothesis of the existence of the set of functions from a set A to a set B , as adequately strong to do most of ordinary mathematics yet still sufficiently constructive.

Later on, Aczel ([1], see also [2] and [3]) picked up the thread and re-cast CST in a more conventional framework, christening this revamped version CZF (Constructive ZF). He also provided a constructive justification for the enterprise, by interpreting CZF in Martin-Löf type theory ([5], [6]). In so doing, he uncovered a principle apparently stronger than Exponentiation, Subset Collection. Subset Collection is most easily understood via an equivalent, Fullness. A set C of total relations from A to B is said to be full if every total relation from A to B can be thinned to one in C . The Fullness Axiom says that every pair of sets has a full set of relations. This principle is valid in Aczel's interpretation, and so is included in CZF.

This all leads to asking what the exact relationship among these axioms actually is. It is easily seen that Power Set implies Subset Collection, which itself implies Exponentiation. What about the reverse implications? SC was chosen over Exponentiation because it's presumably stronger, and Power Set not because it's presumably not Martin-Löf interpretable. Presumably. In this paper, we show that the presumed non-implications are indeed correct. In the next section, it is show that SC does not imply Power Set, and in the one after that, that Exponentiation does not imply SC. In the final section, we observe that a result about Reflection is a by-product of the earlier work.

In fact, although these questions came up in the context of CZF, slightly more than independence over the remaining CZF axioms is shown. The other difference between CZF and IZF is that CZF contains Restricted Separation (i.e. Separation for Δ_0 definable properties), whereas IZF has full Separation. The independence results below include full Separation, and so are actually independence results over IZF - Power Set.

For the record, the axioms of IZF are: Empty, Infinity, Pair, Union, Extensionality, Set Induction, Power Set, Separation, and Collection. CZF consists of Empty, Strong Infinity, Pair, Union, Extensionality, Set Induction, Subset Collection, Δ_0 Separation, and Strong Collection. Strong Infinity is a variant of Infinity that more directly produces ω without resort to Power Set or full Separation. Similarly, Strong Collection is a variant of Collection useful in the absence of full Separation. Collection asserts the existence of a bounding set (for a total function on a set defined by a formula). In the presence of Separation, one can take those elements of a postulated bounding set that come from something in the domain. Strong Collection is the assertion that there is a bounding set consisting only of range values for something in the domain.

2 Subset Collection does not imply Power Set

We will provide a Kripke model of CZF + \neg Power Set (within the meta-theory ZF). Furthermore, full Separation will hold.

Given a partial order P , the full Kripke model \mathbf{M} on P is defined inductively on $\alpha \in \text{ORD}$ as follows. At node σ , the universe \mathbf{M}_σ^α consists of those functions f with domain $P_{\geq \sigma}$ such that for $\tau \geq \sigma$ $f(\tau) \subseteq \bigcup_{\beta < \alpha} \mathbf{M}_\tau^\beta$ and for $\rho \geq \tau$ $f_{\tau\rho}''f(\tau) \subseteq f(\rho)$ (where $f_{\tau\rho}$ is the transition function from node τ to node ρ). $f_{\sigma\tau}$ is then extended from $\bigcup_{\beta < \alpha} \mathbf{M}_\sigma^\beta$ to \mathbf{M}_σ^α as restriction: $f_{\sigma\tau}(f) = f \upharpoonright P_{\geq \tau}$. The universe at σ , \mathbf{M}_σ , is defined as $\bigcup_{\alpha \in \text{ORD}} \mathbf{M}_\sigma^\alpha$. Then \in and $=$ are as easy as pie to define: $\sigma \vdash f \in g$ iff $f \in g(\sigma)$, and $\sigma \vdash f = g$ if $f = g$.

Equivalently, one could simplify the inductive generation of the Kripke sets and the transition functions, at the expense of \in and $=$, by having the Kripke sets have domain all of P : $\mathbf{M}^\alpha = \{f : P \rightarrow \bigcup_{\beta < \alpha} \mathbf{M}^\beta \mid f \text{ is monotonically non-decreasing (i.e. } \sigma < \tau \rightarrow f(\sigma) \subseteq f(\tau))\}$. Then \mathbf{M}_σ is just $\bigcup_{\alpha \in \text{ORD}} \mathbf{M}^\alpha$ (which, note, is independent of σ), and the transition functions are just the

identity. The complication pops up again when defining \in and $=$, since at σ all nodes not extending σ are unimportant but still in the domains and hence need to be actively disregarded: $\sigma \vdash f \in g$ if $\exists h \in g(\sigma) \sigma \vdash f = h$, and $\sigma \vdash f = g$ if $\forall \tau \geq \sigma [\forall h \in f(\tau) \tau \vdash h \in g(\tau) \wedge \forall h \in g(\tau) \tau \vdash h \in f(\tau)]$.

The model we ultimately want to consider differs from that above in that the Kripke sets are to be, as functions, eventually constant. There are several ways this could be achieved. The formalism developed below will be to have P be a class; in particular, ORD. Since every Kripke set is a set, it can give information about only a finite initial segment of P , after which it is assumed to be constant. Alternatively, P could be taken to be ω and the Kripke sets to be eventually hereditarily constant. In fact, this latter approach could be read off from the former by taking a cofinal ω -sequence through ORD (from outside of \mathbf{V} , naturally) and cutting the full model down to those nodes. (That this works might be construed as a reminder that ω is still the largest large cardinal of all.) Neither approach seems essentially easier than the other, so, perhaps arbitrarily, we will henceforth deal only with the former.

At this point it would be prudent, instead of formalizing the ideas above in their full generality, to particularize to the partial order of interest. So let P be ORD. The nodes in P will typically be referred to with letters from the middle of the Greek alphabet ($\kappa, \lambda, \mu, \nu$ and so on) so as not to cause confusion with the use of α, β , and γ used in inductions.

Both constructions above can be adapted to this choice of P . The adaptation of the second construction will be described informally, followed by a thorough description using the first construction, which we remain the basis of this paper. So inductively a Kripke set is a function from $P = \text{ORD}$ to the class of previously defined Kripke sets, which is not only monotonic but also eventually constant: $\mathbf{M}^\alpha = \{f : \text{ORD} \rightarrow \bigcup_{\beta < \alpha} \mathbf{M}^\beta \mid f \text{ is monotonically non-decreasing and } \exists \kappa \forall \lambda > \kappa f(\lambda) = f(\kappa)\}$. Since $f(\kappa)$ is set-sized, not only is each of its elements eventually constant (inductively), but there is a uniform point of stability: $\exists \kappa \forall g \in TC(f)$ (transitive closure) g is constant beyond κ . The transition functions, as well as membership and equality, as are above.

To formalize this as stated, an axiomatic theory of classes would be necessary. To work within a more conventional framework, meaning a theory of sets (the framework of this paper being, after all, ZF), an element f will have to be a function on a (necessarily) set-sized initial segment of ORD, beyond which f is implicitly constant. The technical work here is to define the transition functions $f_{\kappa\lambda}$ on f , even when $\lambda \notin \text{dom}(f)$. Since this involves mucking with the domains anyway, it's easier to adapt the former of the constructions above. The reader with a clear intuition who is willing to trust that these matters can be formalized may safely skip to the proposition.

The full Kripke model \mathbf{M} on P is defined inductively on $\alpha \in \text{ORD}$ as follows. At node κ , the universe \mathbf{M}_κ^α consists of those functions f with domain $\lambda \setminus \kappa$ (for some $\lambda > \kappa$) such that for $\mu \in \text{dom}(f)$ $f(\mu) \subseteq \bigcup_{\beta < \alpha} \mathbf{M}_\mu^\beta$ and for $\nu \in \text{dom}(f)$, $\nu \geq \mu$, $f_{\mu\nu}''f(\mu) \subseteq f(\nu)$ (where $f_{\mu\nu}$ once again is the transition function from

node μ to node ν). Without loss of generality we can also require that for all $\mu \in \text{dom}(f)$ and $g \in f(\mu)$ $\text{dom}(f) \supseteq \text{dom}(g)$. (“WLOG” means here that this additional stipulation does not exclude any sets. More precisely, if f satisfies the definition for a set in the Kripke structure except for the condition on the domain, then there is a Kripke set g satisfying the domain condition such that $\kappa \models “f=g”$. Such a g would be constructed as an end-extension of f , whereby for $\mu \in \text{dom}(g) \setminus \text{dom}(f)$, $g(\mu)$ would be taken as $f_{\kappa\mu}(f)$, using the transition functions about to be defined.)

$f_{\kappa\mu}$ is then extended from $\bigcup_{\beta < \alpha} \mathbf{M}_{\kappa}^{\beta}$ to $\mathbf{M}_{\kappa}^{\alpha}$ as follows. If $\mu \in \text{dom}(f)$, then $f_{\kappa\mu}$ is restriction: $f_{\kappa\mu}(f) = f \upharpoonright \text{dom}(f) \setminus \mu$. If $\mu \notin \text{dom}(f)$, then $f_{\kappa\mu}(f)$ will have domain $\{\mu\}$ (in the notation from the last paragraph, λ will be $\mu + 1$), and $(f_{\kappa\mu}(f))(\mu)$ will be (working inductively here) $\bigcup_{\nu \in \text{dom}(f)} f_{\nu\mu} f(\nu)$. It is left to the reader to show that the transition functions compose as they are supposed to (i.e. $f_{\lambda\mu} \circ f_{\kappa\lambda} = f_{\kappa\mu}$).

The universe at node κ , \mathbf{M}_{κ} , can now be defined as $\bigcup_{\alpha \in \text{ORD}} \mathbf{M}_{\kappa}^{\alpha}$. The transition functions $f_{\kappa\mu}$ defined along the way on $\mathbf{M}_{\kappa}^{\alpha}$ cohered, which is why the dependence on α was dropped from the notation; the ultimate $f_{\kappa\mu}$ defined on \mathbf{M}_{κ} is the union of the partial $f_{\kappa\mu}$ ’s defined along the way.

We will need that the model \mathbf{M} is definable in \mathbf{V} , so we may as well check each building block in the development of \mathbf{M} when it’s first introduced. Notice that $\mathbf{M}_{\kappa}^{\alpha}$ cannot be a function of α , since $\mathbf{M}_{\kappa}^{\alpha}$ is not even a set! However, the relation “ $f \in \mathbf{M}_{\kappa}^{\alpha}$ ” (as a relation on f , α , and κ), as well as “ $f_{\kappa\mu}(f) = g$ ” (as a relation on f , g , κ , and μ) is Δ_1 . This can be shown via a simultaneous Σ_1 induction. The central point in this argument is that the definitions (as given above) contain only bounded quantification and recursive calls.

\in and $=$ must then be defined via a mutual induction to make sure that Extensionality holds in the end. The problem here, which does not come up in the case of a set-sized underlying p.o., is illustrated by the fact that both $\{\langle 0, \emptyset \rangle\}$ and $\{\langle 0, \emptyset \rangle, \langle 1, \emptyset \rangle\}$ stand for the empty set yet are not themselves equal as sets. The reader uninterested in such details can safely skip to the main proposition.

Assume that \in_{β} and $=_{\beta}$ are defined on $\mathbf{M}_{\kappa}^{\beta}$ for all nodes κ and, inductively, for all stages $\beta < \alpha$. For $f, g \in \mathbf{M}_{\kappa}^{\alpha}$, $\kappa \models “f \in_{\alpha} g”$ iff for some $h \in g(\kappa)$ and $\beta < \alpha$, $\kappa \models “f =_{\beta} h”$. Notice that the sequence of \in_{α} s is monotonically non-decreasing, that is, if $f \in_{\alpha} g$ and $\gamma > \alpha$ then $f \in_{\gamma} g$. This holds because the defining condition of \in_{α} remains true if α is increased. Regarding $=$, $\kappa \models “f =_{\alpha} g”$ iff for all $\lambda \in \text{dom}(f) \cup \text{dom}(g)$, $h \in (f_{\kappa\lambda}(f))(\lambda)$, $\lambda \models “h \in_{\alpha} f_{\kappa\lambda}(g)”$, and for all $h \in (f_{\kappa\lambda}(g))(\lambda)$, $\lambda \models “h \in_{\alpha} f_{\kappa\lambda}(f)”$. Notice that the sequence of $=_{\alpha}$ s is non-decreasing, because, in the defining condition, α appears only in \in_{α} , \in_{α} appears only positively, and \in_{α} is already known to be non-decreasing. Furthermore, \in_{α} and $=_{\alpha}$ are non-increasing for a fixed collection of sets. That is, for $\alpha = \max(\text{rk } f, \text{rk } g)$, $\kappa \models “f \in_{\alpha} g”$ iff for some $\gamma > \alpha$ $\kappa \models “f \in_{\gamma} g”$, and similarly for $=_{\alpha}$. This can be proved by a mutual induction. Finally, let \in be

$\bigcup_{\beta \in ORD} \in_\beta$, i.e. $\kappa \models "f \in g"$ iff for some β $\kappa \models "f \in_\beta g"$, and similarly for $=$.

Considering definability again, both \in and $=$ are Δ_1 -definable ternary relations (i.e. $\kappa \models f \in g$ and $\kappa \models f = g$). This is based on two facts. First, for $\alpha = \max(\text{rk } f, \text{rk } g)$, $\kappa \models "f \in g"$ iff $\kappa \models "f \in_\alpha g"$, and similarly for $=$. Second, \in_α and $=_\alpha$ are defined by a simultaneous recursion which contains only bounded quantification and recursive calls.

Lemma 2.0.1 *The transition functions respect the primitive relations \in and $=$.*

proof: It suffices to show that the transition functions respect \in_α and $=_\alpha$. This can be done by a simultaneous induction on α . The case of \in_α is trivial. To show $\kappa \models f =_\alpha g$ implies $\mu \models f_{\kappa\mu} =_\alpha f_{\kappa\mu}(g)$, if $\mu \in \text{dom}(f) \cup \text{dom}(g)$, then this is trivial, since the only effect of $f_{\kappa\mu}$ is to shrink the set of λ 's that need be considered. Otherwise one has to unravel the definition of $f_{\kappa\mu}$, which again brings us to consider those λ s in $\text{dom}(f) \cup \text{dom}(g)$. The details are left to the reader. ■

Lemma 2.0.2 *$=$ (as defined above) satisfies the equality axioms.*

proof: 1. $0 \models \forall x \, x = x$: Assume inductively that for all $x \in \mathbf{M}_\kappa^\beta$ ($\beta < \alpha$) $\kappa \models x =_\beta x$ and that $f \in \mathbf{M}_\kappa^\alpha$. We need to show that $\kappa \models f =_\alpha f$, i.e. for all $\lambda \in \text{dom}(f)$, $h \in (f_{\kappa\lambda}(f))(\lambda)$, $\lambda \models "h \in_\alpha f_{\kappa\lambda}(f)"$. Since $\lambda \in \text{dom}(f)$, $f_{\kappa\lambda}$ operates via restriction, so the condition to be proved simplifies to for all $\lambda \in \text{dom}(f)$, $h \in f(\lambda)$, $\lambda \models "h \in_\alpha f_{\kappa\lambda}(f)"$. Given such a λ and h , by the definition of \in_α , we need a $g \in f(\lambda)$ and β so that $\lambda \models "h =_\beta g"$. Letting g be h and β be $\text{rk}(h)$, the inductive hypothesis gives exactly this.

2. $0 \models \forall x, y \, x = y \rightarrow y = x$: Trivial, because the definition of \in_α was symmetric in f and g .

3. $0 \models \forall x, y, z \, x = y \wedge y = z \rightarrow x = z$: As in 1, the transitivity of $=_\alpha$ is proven inductively on α . The technical difficulty here is that we must consider the domains of arbitrary Kripke sets f , g , and h , which might all differ. Possibly the easiest way to do this is to start with the observation that, for arbitrary $\lambda \geq \kappa$, $(f_{\kappa\lambda}(f))(\lambda) = \bigcup_{\mu \leq \lambda, \mu \in \text{dom}(f)} f_{\mu\lambda} f(\mu)$. So when evaluating $\kappa \models f =_\alpha h$, after being led to consider $i \in (f_{\kappa\lambda}(f))(\lambda)$, we can conclude that $i \in f_{\mu\lambda} f(\mu)$ for some $\mu \in \text{dom}(f)$. Using a pre-image i_μ of i in $f(\mu)$ and the postulated equality of f and g , one finds a $\nu \in \text{dom}(g)$, $\nu \leq \mu$, a pre-image i_ν of i_μ , and a $j_\nu \in g(\nu)$ such that $i_\nu =_\beta j_\nu$. Then one finds a $\xi \leq \nu$, $j_\xi \in g(\xi)$, $k_\xi \in h(\xi)$ such that $j_\xi =_\gamma k_\xi$. At this point one uses the monotonicity of $=$ to move up to $\delta = \max(\beta, \gamma)$, the inductively assumed transitivity of $=_\delta$ to get that $i_\xi =_\delta k_\xi$, and the fact that the transition functions respect the primitive relations to conclude $\lambda \models "i \in_\alpha h"$. The details are left to the reader.

4. $0 \models \forall x, y, z \ x = y \wedge x \in z \rightarrow y \in z$: If $\kappa \models "f=g \wedge f \in h"$, then $\kappa \models "f=i"$ for some $i \in h(\kappa)$. By the transitivity of $=$, $\kappa \models "g=i"$.
5. $0 \models \forall x, y, z \ x = y \wedge z \in x \rightarrow z \in y$: If $\kappa \models "f=g \wedge h \in f"$, then for some $i \in f(\kappa)$ $\kappa \models "h=i"$, and also $\kappa \models "i \in g"$. That means that for some $j \in g(\kappa)$ $\kappa \models "i=j"$. By the transitivity of $=$, $\kappa \models "h=j"$, and hence also $"h \in g"$.

With \in and $=$ defined, the description of the model \mathbf{M} is complete. The definability of \mathbf{M} follows automatically from the work above. Since the universes at each node and the primitive relations (there being no primitive functions) are definable, the definability of the satisfaction relation $\kappa \models \phi$ comes directly from the inductive determination of satisfaction in Kripke models. Furthermore, while we make no claim along the lines of Σ_n -satisfaction being Σ_n definable, it is true that Δ_0 -satisfaction (i.e. $\kappa \models \phi$, ϕ a Δ_0 formula) is Δ_1 definable, as follows. In the inductive definition of satisfaction in Kripke models (when restricting to Δ_0 formulas), the only times an unbounded quantifier is introduced are for the clauses \forall and \rightarrow , when the collection of nodes is quantified over. Such quantification can be eliminated by restricting it to the maximum κ of the domains of the parameters, because the formula is true at κ iff it's true at any later node, this last fact being provable inductively. (A different proof, which shows the same fact without the restriction of ϕ being Δ_0 , is that there is an isomorphism between the Kripke model from κ on and the model from λ on, where λ is any ordinal $> \kappa$. This latter argument is given in detail in the proof of Separation later in this section.)

What remains is to show that \mathbf{M} has the desired properties.

Theorem 2.0.3 $\mathbf{M} \models \text{CZF} + \neg \text{Power Set}$

proof: The simpler axioms of CZF are Pairing, Union, and Infinity, and the proofs that they hold in \mathbf{M} are left to the reader. What we will show are Extensionality, Set Induction, Restricted Separation, Subset Collection, Strong Collection, and, of course \neg Power Set. As promised above, we will go beyond CZF and show full Separation.

Extensionality: What needs to be shown is $\forall x \forall y (x = y \leftrightarrow \forall z (z \in x \leftrightarrow z \in y))$. The implication \rightarrow follows easily from equality property 5. in the previous proposition. For the implication \leftarrow , given sets f and g and node κ , let $\alpha = \max(\text{rk } f, \text{rk } g)$. To show that $f =_\alpha g$ is precisely, by definition, $\forall \lambda \in \text{dom}(f) \cup \text{dom}(g), h \in (f_{\kappa\lambda}(f))(\lambda), \lambda \models "h \in_\alpha f_{\kappa\lambda}(g)"$, and vice versa. So let h be a member of $(f_{\kappa\lambda}(f))(\lambda)$. By assumption, we have $\kappa \models "\forall z (z \in f \leftrightarrow z \in g)"$. This holds in particular for h : $\lambda \models "h \in (f_{\kappa\lambda}(f)) \leftrightarrow h \in (f_{\kappa\lambda}(g))"$. In order to use this, we need that $\lambda \models "h \in (f_{\kappa\lambda}(f))"$, when what we have is merely $h \in (f_{\kappa\lambda}(f))(\lambda)$. Unraveling the definition of $\lambda \models x \in_\beta y$ and making the obvious choices along the way, this reduces to showing that $\lambda \models h =_\beta h$ for some β . Letting β be $\text{rk}(h)$,

this follows from the proof (albeit not the statement) of the first equality axiom (previous proposition). Hence we can conclude that $\lambda \models "h \in (f_{\kappa\lambda}(f)) \leftrightarrow h \in (f_{\kappa\lambda}(g))"$. What we still need is that $\lambda \models "h \in_{\alpha} f_{\kappa\lambda}(g)"$. That this is the case was already observed while defining \in_{α} and $=_{\alpha}$. The other direction ("vice versa") is symmetric in f and g . The details are left to the reader.

Set Induction: This is easy, because the Kripke sets were defined inductively. That is, given an appropriate property $\phi(x)$, one would show that $\phi(x)$ holds for all members of \mathbf{M}^{α} inductively on α .

Restricted Separation: In fact, full Separation holds. Let κ be a node, a a set in \mathbf{M}_{κ} , and ϕ a formula (with free variable x) over (i.e. with parameters from) \mathbf{M}_{κ} . Let $\lambda > \kappa$ include not only $\text{dom}(a)$ but also the domains of all of ϕ 's parameters. For $\mu > \kappa$ we will use the notation $f_{\kappa\mu}(\phi)(x)$ to mean $\phi(x, f_{\kappa\mu}(b_0), \dots, f_{\kappa\mu}(b_{n-1}))$, where the b_i 's list all of ϕ 's parameters. In \mathbf{V} , let $\text{Sep}_{a,\phi}$ be a function with domain $\lambda + 1 \setminus \kappa$ such that for all μ in this domain

$$\text{Sep}_{a,\phi}(\mu) = \{x \in (f_{\kappa\mu}(a))(\mu) \mid \mu \models f_{\kappa\mu}(\phi)(x)\}.$$

This function $\text{Sep}_{a,\phi}$ is definable in \mathbf{V} because \mathbf{M} is, and it should be clear that for all $\mu \in \text{dom}(\text{Sep}_{a,\phi})$, $x \in \mathbf{M}_{\mu}$,

$$\mu \models "x \in f_{\kappa\mu}(\text{Sep}_{a,\phi}) \leftrightarrow x \in f_{\kappa\mu}(a) \wedge f_{\kappa\mu}(\phi)(x)".$$

We need only show why the same holds also for $\mu > \lambda$. The idea is that beyond λ nothing relevant (i.e. a or ϕ) is changing anymore, and so $\text{Sep}_{a,\phi}$ need not change either. In detail, for any such μ , by the choice of λ , $f_{\lambda\mu}$ is an isomorphism on $f_{\kappa\lambda}(a)$, $f_{\kappa\lambda}(\phi)$, and $f_{\kappa\lambda}(\text{Sep}_{a,\phi})$, the nature of which is replacing the domains in the transitive closures of the sets in question, which are all $\{\lambda\}$, with $\{\mu\}$. This isomorphism can be extended to one between \mathbf{M}_{λ} and \mathbf{M}_{μ} by hereditarily replacing ordinals in the domains, necessarily of the form $\lambda + \xi$, with $\mu + \xi$. Under this isomorphism, $\{x \in \mathbf{M}_{\lambda} \mid \lambda \models "x \in f_{\kappa\lambda}(a) \wedge f_{\kappa\lambda}(\phi)(x)"\}$, i.e. $\text{Sep}_{a,\phi}(\lambda)$, gets mapped to $\{x \in \mathbf{M}_{\mu} \mid \mu \models "x \in f_{\lambda\mu}(f_{\kappa\lambda}(a)) \wedge f_{\lambda\mu}(f_{\kappa\lambda}(\phi))(x)"\}$. By the composition property of transition functions ($f_{\sigma\tau} \circ f_{\rho\sigma} = f_{\rho\tau}$, a necessary condition on all Kripke models), we have the desired result.

Subset Collection: Let $\kappa \in \text{ORD}$, $a, b \in \mathbf{M}_{\kappa}$. Let $\lambda \supseteq \text{dom}(a)$, $\text{dom}(b)$. $\text{SubColl}_{a,b}$ will have domain $\lambda + 1 \setminus \kappa$. For all μ in this domain, let

$$\text{SubColl}_{a,b}(\mu) = \{x \in \mathbf{M}_{\mu} \mid \mu \models "x \text{ is a total relation from } a \text{ to } b"\}.$$

It is immediate that for all $\mu \in \text{dom}(\text{SubColl}_{a,b})$, $x \in \mathbf{M}_{\mu}$,

$$\mu \models "x \in f_{\kappa\mu}(\text{SubColl}_{a,b}) \leftrightarrow x \text{ is a total relation from } a \text{ to } b".$$

For $\mu > \lambda$, the idea is that from λ to μ neither a nor b has changed. So if node μ forces x to be a total relation from a to b , then any ordered pairs that might enter x at a later node can be disregarded and x would still remain total, since a doesn't grow any more. Furthermore, the ordered pairs at node μ can be

“pulled back” to node λ , since neither a nor b grew from λ to μ . This (possibly truncated) version of x was put into $\text{SubColl}_{a,b}$ by at latest node λ , and is indeed a total sub-relation of x .

In more detail, observe that for any such μ , by the choice of λ , $f_{\lambda\mu}$ is an isomorphism on $f_{\kappa\lambda}(a)$ and $f_{\kappa\lambda}(b)$. Let $\mu \models “x \text{ is a total relation from } f_{\kappa\mu}(a) \text{ to } f_{\kappa\mu}(b)”$. Then $x(\mu)$ is a total relation from $(f_{\kappa\mu}(a))(\mu)$ to $(f_{\kappa\mu}(b))(\mu)$. This induces via the isomorphism a total relation from $(f_{\kappa\lambda}(a))(\lambda)$ to $(f_{\kappa\lambda}(b))(\lambda)$; let z be the function with domain $\{\lambda\}$ and $z(\lambda)$ being this relation. $\lambda \models “z \text{ is total}”$ since a doesn’t change beyond λ , so $\lambda \models “z \in \text{SubColl}_{a,b}”$. By the construction of z , $\mu \models f_{\lambda\mu}(z) \subseteq x$.

Strong Collection: Let κ be a node, a a set in \mathbf{M}_κ , and ϕ a formula (with free variable x) over (i.e. with parameters from) \mathbf{M}_κ , such that $\kappa \models “\forall x \in a \exists y \phi(x, y)”$. Let $\lambda > \kappa$ include not only $\text{dom}(a)$ but also the domains of all of ϕ ’s parameters. As above, the notation $f_{\kappa\mu}(\phi)(x)$ means $\phi(x, f_{\kappa\mu}(b_0), \dots, f_{\kappa\mu}(b_{n-1}))$, where the b_i ’s list all of ϕ ’s parameters. In \mathbf{V} , a function $\text{StrColl}_{a,\phi}$ with domain $\lambda + 1 \setminus \kappa$ can be constructed so that for all μ and ν in this domain with $\mu < \nu$,

$$\forall x \in a(\mu) \exists y \in \text{StrColl}_{a,\phi}(\mu) \mu \models \phi(x, y),$$

$$\forall y \in \text{StrColl}_{a,\phi}(\mu) \exists x \in a(\mu) \mu \models \phi(x, y),$$

and $f_{\mu\nu}”\text{StrColl}_{a,\phi}(\mu) \subseteq \text{StrColl}_{a,\phi}(\nu)$. $\text{StrColl}_{a,\phi}$ clearly acts as a bounding set at all nodes through λ . Beyond λ the idea is that nothing relevant (i.e. a or ϕ) is changing anymore, and so $\text{StrColl}_{a,\phi}$ need not change either. The argument is essentially the same as that for Restricted Separation, and is left to the reader.

\neg **Power Set:** For $\kappa \in \text{ORD}$, let 1_κ be the Kripke set with domain $\kappa + 1$ such that for $\lambda < \kappa$ $1_\kappa(\lambda) = \emptyset$ and $1_\kappa(\kappa) = \langle \kappa, \emptyset \rangle$. $0 \models “1_\kappa \subseteq 1 =_{\text{def}} \{0\}”$, but by the set-sized nature of the members of \mathbf{M} , no set contains all of the 1_κ ’s. Therefore 1 has no power set in \mathbf{M} . \blacksquare

3 Exponentiation does not imply Subset Collection

We will provide a Kripke model of CZF - Subset Collection + Exponentiation + \neg Subset Collection. Furthermore, full Separation will hold.

The model to be built in this section is an extension of that from the previous section, and ultimately can be defined only in a generic extension of \mathbf{V} . These sets, though, will be named by terms in \mathbf{V} and so we will work in \mathbf{V} at first to identify the terms. When the terms are interpreted, they will be functions just as in the last section, and so their domains will also be given during their inductive generation.

The terms are defined inductively on the ordinals. At stage 0, at any node κ , there is a term $f_{0\kappa}(G_\lambda)$, which will be interpreted as a function with domain $[\kappa, \max(\kappa, \lambda)]$. (For $\kappa = 0$, these will be referred to more simply as G_λ .) (Of course, in this context the symbol “ $f_{0\kappa}$ ” is just a piece of syntax. It should, however, remind you of the transition function from node 0 to κ . In an attempt to reduce notation, we will use the same notation for a term as for its interpretation, hence the double meaning of “ $f_{0\kappa}$ ”.) At any later stage, if a is a term at node κ , then so is $\text{Union}(a)$ (with the same domain); if a is a term at node κ , ϕ is a Δ_0 formula (in which all of the parameters are terms at node κ), and $\lambda > \kappa$, then $\text{Sep}_{a,\phi,\lambda}$ is a term at node κ with domain $[\kappa, \lambda]$; and if a function is build via the inductive construction of the previous section (given at the bottom of p. 3), then it is also a term in which the node and domain of definition are explicit.

Concerning the interpretation of these terms, we introduce some notation. For an integer n and node κ , $n_\kappa^{\mathbf{M}}$ is the internal version of n ; that is, $0_\kappa^{\mathbf{M}} = \langle \kappa, \emptyset \rangle$ and $n+1_\kappa^{\mathbf{M}} = \langle \kappa, n_\kappa^{\mathbf{M}} \cup \{n_\kappa^{\mathbf{M}}\} \rangle$. Let $\mathbf{N}_\kappa^{\mathbf{M}}$ be $\{n_\kappa^{\mathbf{M}} \mid n \in \mathbf{N}\}$. (Notice that $\mathbf{N}_\kappa^{\mathbf{M}}$ is not actually a set \mathbf{M} ; if we had instead wanted to define the internal version of the natural numbers, say, that would be $\langle \kappa, \mathbf{N}_\kappa^{\mathbf{M}} \rangle$ (or a variant thereof).)

Let G be a generic relation on \mathbf{N} . That is, G is generic over \mathbf{V} for the forcing partial order which consists of functions from finite subsets of $\mathbf{N} \times \mathbf{N}$ to 2; p extends q (written “ $p \geq q$ ”) if p extends q as a function, i.e. $p \supseteq q$. A generic can be identified with a total function $F : \mathbf{N} \times \mathbf{N} \rightarrow 2$; letting G be $\{ \langle m, n \rangle \mid F(\langle m, n \rangle) = 1 \}$, G is then a generic relation on \mathbf{N} . Notice that by genericity, both G and G^{-1} are total.

In $\mathbf{V}[G]$, let $G_\kappa^{\mathbf{M}}$ be $\{ \langle m_\kappa^{\mathbf{M}}, n_\kappa^{\mathbf{M}} \rangle^{\mathbf{M}} \mid \langle m, n \rangle \in G \}$ (where $\langle \cdot, \cdot \rangle^{\mathbf{M}}$ means the \mathbf{M} -internal ordered pair function). Similarly, let $(\mathbf{N} \times \mathbf{N})_\kappa^{\mathbf{M}}$ be $\{ \langle m_\kappa^{\mathbf{M}}, n_\kappa^{\mathbf{M}} \rangle^{\mathbf{M}} \mid m, n \in \mathbf{N}_\kappa^{\mathbf{M}} \}$.

Using the $G_\kappa^{\mathbf{M}}$ s, we can now provide an interpretation for the terms. (Remember, we are using the same notation for the interpretation of a term as for the term itself; whether a term or its interpretation is meant should be clear from the context.) For $f_{0\kappa}(G_\lambda)$, $f_{0\kappa}(G_\lambda)(\mu) = G_\mu^{\mathbf{M}}$ for $\mu < \lambda$ (the internal version of G), and $f_{0\kappa}(G_\lambda)(\lambda) = (\mathbf{N} \times \mathbf{N})_\kappa^{\mathbf{M}}$ otherwise (the full relation on \mathbf{N}). $\text{Union}(a)(\kappa)$ (where $\kappa \in \text{dom}(a)$) is $\bigcup_{f \in a(\kappa)} f(\kappa)$. The interpretation of $\text{Sep}_{a,\phi,\lambda}$ is similar to that of $\text{Sep}_{a,\phi}$ in the previous section (See the proof of Restricted Separation), only there an essentially arbitrary λ was chosen (the only requirement being that λ be big enough), whereas here the λ in question is a part of the term and the size restriction is dropped. So for $\mu \in [\kappa, \lambda]$,

$$\text{Sep}_{a,\phi,\lambda}(\mu) = \{x \in (f_{\kappa\mu}(a))(\mu) \mid \mu \models f_{\kappa\mu}(\phi)(x)\}.$$

The interpretation of a function is straightforward: mod the above mentioned identification of terms with their interpretations, there is literally nothing to do.

Of course, to have a Kripke model, we need some other facts and definitions; indeed, the interpretation above is dependent upon some of these, which would therefore need to be proven and defined simultaneously with this inductive in-

terpretation. That is, we need to define transition functions $f_{\kappa\mu}$, show that they cohere (i.e. $f_{\lambda\mu} \circ f_{\kappa\lambda} = f_{\kappa\mu}$), and show that the Kripke sets are non-decreasing (i.e. if $\kappa \models a \in b$ then $\mu \models f_{\kappa\mu}(a) \in f_{\kappa\mu}(b)$). We will leave the details to the reader, and content ourselves with one observation. Since λ was pointedly left out of the domain of $\text{Sep}_{a,\phi,\lambda}$, for all $\mu \geq \lambda$ including λ itself

$$\text{Sep}_{a,\phi,\lambda}(\mu) = \bigcup_{\nu \in [\kappa, \text{lambda})} f_{\nu\mu}'' \text{Sep}_{a,\phi,\lambda}(\nu).$$

Notice that \mathbf{M} (i.e. $\kappa \models \phi$) is definable in $\mathbf{V}[\mathbf{G}]$. In particular, as in the previous section, Δ_0 -satisfaction is Δ_1 . Furthermore, since \mathbf{G} is set-generic over \mathbf{V} , by general forcing technology, $p \Vdash \text{“}\kappa \models \phi\text{”}$ is definable in \mathbf{V} , and, in particular, $p \Vdash \text{“}\kappa \models \phi\text{”}$, $\phi \Delta_0$, is Δ_1 .

Theorem 3.0.4 $\mathbf{M} \models \text{CZF} - \text{Subset Collection} + \text{Exponentiation} + \neg \text{Subset Collection}$

proof: Extensionality, Set Induction, and Infinity are left to the reader.

Union: This follows from the presence (and meaning) of the terms of the form $\text{Union}(a)$.

Pairing: Given $a, b \in \mathbf{M}_\kappa$, let λ include $\text{dom}(a)$ and $\text{dom}(b)$. For $\mu \in \lambda \setminus \kappa$ let $f(\mu)$ be $\{f_{\kappa\mu}(a), f_{\kappa\mu}(b)\}$, which suffices. Notice that f is definable even if a or b is just a term, so that f exists at least as a term in \mathbf{V} (even if inductively uninterpretable in \mathbf{V}).

Restricted Separation: This follows from the presence (and meaning) of the terms of the form $\text{Sep}_{a,\phi,\lambda}$, with the proof being essentially that in the previous section. The one change is that the isomorphism given there must be extended to account for terms of the form $\mathbf{G}\kappa$ (the other terms already being accounted for inductively). Under this isomorphism, $\mathbf{G}_{\lambda+\xi}$ is to be sent to $\mathbf{G}_{\mu+\xi}$.

In fact, full Separation holds, even though the separation terms are only for Δ_0 formulas. The first guess as to how to proceed is as follows. Given a and ϕ a term and formula in (resp. over) \mathbf{M}_κ , let $\lambda > \kappa$ include the domains of a and of ϕ 's parameters. Let $\mu > \lambda$ be such that $\mathbf{V}_\mu[\mathbf{G}]$ contains a , ϕ 's parameters, and is a Σ_n elementary substructure of $\mathbf{V}[\mathbf{G}]$ for n sufficiently large (i.e. large enough to reflect the truth of $\phi(x)$ in \mathbf{M}). Let f be the function with domain $\mu \setminus \kappa$ such that $f(\nu) = \{g \in \mathbf{V}_\mu \mid g \text{ is a term naming an element of } \mathbf{M}_\nu\}$. The notation “ ϕ^f ” means “ ϕ with all of the unbounded variables bounded by f ”. By the choice of μ , $\phi(x)$ is interpreted in \mathbf{M} (in $\mathbf{V}[\mathbf{G}]$) just as $\phi(x)$ would be in $\mathbf{V}_\mu[\mathbf{G}]$'s version of \mathbf{M} . If in turn that could be captured as $\text{Sep}_{a,\phi^f,\mu}$ (in $\mathbf{V}[\mathbf{G}]$), then we'd be done.

The problem is that, while f reflects the restriction to \mathbf{M}^μ , it cannot reflect the restriction of P from ORD to μ . To give a concrete example, let ϕ be the sentence “ $\forall x \forall y x = y \vee x \neq y$ ”. In \mathbf{M} , this sentence is never true, and so any separation set for ϕ should be the empty set. However, for any choice of f , ϕ^f

eventually becomes true (at latest beyond f 's domain), so $\text{Sep}_{a,\phi^f,\mu}$ eventually becomes a .

We need some way of saying “don't look beyond this node.” This can be done with a technique developed in its essence by Friedman and Scedrov. In [4], they consider a particular formula ψ (in the notation to be introduced) and a particular Kripke model; moreover, their model was a two-node model, so their $\mathbf{M}_{pre-\psi}$ (see below) was a classical and not even a Kripke structure. Nonetheless, their work is easily generalized to all formulas and all models. To have this generalization shown, and to keep this paper self-contained, we develop the details.

Let ψ be a sentence, possibly with parameters from the universe at the bottom node \perp in some Kripke model \mathbf{M} . Using ψ , a transformation on the formulas can be defined inductively. If ϕ is atomic (including \perp , the always false statement), then ϕ^* is $\phi \vee \psi$. The inductive clauses are transparent to this construction: $(\phi_0 \rightarrow \phi_1)^* = (\phi_0^* \rightarrow \phi_1^*)$, $\forall x \phi(x) = \forall x \phi^*(x)$, etc. It is easy to see inductively on ϕ that for all formulas ϕ and all nodes σ with $\sigma \models \psi$, $\sigma \models \phi^*$. Let $\mathbf{M}_{pre-\psi}$ be \mathbf{M} with all the nodes where ψ holds thrown out. That is, $\mathbf{M}_{pre-\psi}$'s underlying partial order is the initial sub-order of \mathbf{M} 's P consisting of $\{\sigma \mid \sigma \not\models \psi\} =_{def} P_{pre-\psi}$ (which is assumed to be non-empty, i.e. $\perp \not\models \psi$). The universe of $\mathbf{M}_{pre-\psi}$ at σ is the same as for \mathbf{M} , the transition functions are the same, the atomic relations and functions have the same interpretation. Because nodes in $\mathbf{M}_{pre-\psi}$ are also nodes in \mathbf{M} , we will distinguish notationally truth in $\mathbf{M}_{pre-\psi}$ from truth in \mathbf{M} with $\models_{\mathbf{M}_{pre-\psi}}$ versus $\models_{\mathbf{M}}$.

Lemma 3.0.5 $\sigma \models_{\mathbf{M}} \phi^*$ iff $\sigma \models_{\mathbf{M}_{pre-\psi}} \phi$.

proof: By induction on ϕ .

- ϕ atomic: $\sigma \models_{\mathbf{M}} \phi^*$ iff $\sigma \models_{\mathbf{M}} \phi \vee \psi$. Since by hypothesis $\sigma \not\models_{\mathbf{M}} \psi$, this is equivalent to $\sigma \models_{\mathbf{M}} \phi$, which is itself equivalent to $\sigma \models_{\mathbf{M}_{pre-\psi}} \phi$, since $\mathbf{M}_{pre-\psi}$ is a restriction of \mathbf{M} and the atomic relations are defined locally. (Although this lemma is very general, the reader can be excused for thinking of models of set theory, and ask whether equality as defined by “same members” really holds. Note we do not claim that Extensionality is preserved! The equality in \mathbf{M} is taken as primitive.)
- $\phi_0 \vee \phi_1$: $\sigma \models_{\mathbf{M}} \phi_0^* \vee \phi_1^*$ iff $\sigma \models_{\mathbf{M}} \phi_0^*$ or $\sigma \models_{\mathbf{M}} \phi_1^*$ iff (inductively) $\sigma \models_{\mathbf{M}_{pre-\psi}} \phi_0$ or $\sigma \models_{\mathbf{M}_{pre-\psi}} \phi_1$ iff $\sigma \models_{\mathbf{M}_{pre-\psi}} \phi_0 \vee \phi_1$.
- $\phi_0 \wedge \phi_1$: Analogous to \vee .
- $\phi_0 \rightarrow \phi_1$: It has already been observed that: for $\tau \in P \setminus P_{pre-\psi}$ $\tau \models_{\mathbf{M}} \phi_0^* \wedge \phi_1^*$.

So for $\sigma \in P_{pre-\psi}$, $\sigma \models_{\mathbf{M}} \phi_0^* \rightarrow \phi_1^*$ reduces to:
for all extensions $\tau \in P_{pre-\psi}$ of σ , if $\tau \models_{\mathbf{M}} \phi_0^*$ then $\tau \models_{\mathbf{M}} \phi_1^*$.
Inductively, this is equivalent to:
for all extensions $\tau \in P_{pre-\psi}$ of σ , if $\tau \models_{\mathbf{M}_{pre-\psi}} \phi_0$ then $\tau \models_{\mathbf{M}_{pre-\psi}} \phi_1$,
i.e. $\sigma \models_{\mathbf{M}_{pre-\psi}} \phi_0 \rightarrow \phi_1$.

- $\forall x \phi(x)$: It has already been observed that:
for $\tau \in P \setminus P_{pre-\psi}$ and $x \in \mathbf{M}_\tau$, $\tau \models_{\mathbf{M}} \phi^*(x)$.
So for $\sigma \in P_{pre-\psi}$, $\sigma \models_{\mathbf{M}} \forall x \phi^*(x)$ reduces to:
for all extensions $\tau \in P_{pre-\psi}$ of σ and $x \in \mathbf{M}_\tau$, $\tau \models_{\mathbf{M}} \phi^*(x)$.
Inductively, this is equivalent to:
for all extensions $\tau \in P_{pre-\psi}$ of σ , and $x \in \mathbf{M}_\tau$, $\tau \models_{\mathbf{M}_{pre-\psi}} \phi(x)$,
i.e. $\sigma \models_{\mathbf{M}_{pre-\psi}} \forall x \phi(x)$.
- $\exists x \phi(x)$: $\sigma \models_{\mathbf{M}} \exists x \phi^*(x)$ iff
for some $x \in \mathbf{M}_\sigma$ $\sigma \models_{\mathbf{M}} \phi^*(x)$ iff (inductively)
for some $x \in \mathbf{M}_\sigma$ $\sigma \models_{\mathbf{M}_{pre-\psi}} \phi(x)$ iff
 $\sigma \models_{\mathbf{M}_{pre-\psi}} \exists x \phi(x)$.

■

To return to the issue at hand, let ψ be the sentence “ $0 \in 1_\mu$ ”, where μ is as above. (Recall that 1_μ is defined so that ψ is true exactly at nodes μ and beyond.) This choice of ψ determines a $*$ -transformation on formulas. Then $\text{Sep}_{a, \phi^{f*}, \mu}$ (equivalently, $\text{Sep}_{a, \phi^{*f}, \mu}$) is as desired. That is, we claim that

$$\kappa \models \forall x (x \in \text{Sep}_{a, \phi^{*f}, \mu} \leftrightarrow (x \in a \wedge \phi(x))).$$

(In what follows, we suppress mention of the transition functions $f_{\kappa\lambda}$, for purposes of readability, preferring instead to write the incorrect “ $\lambda \models x \in a$ ” instead of “ $\lambda \models x \in f_{\kappa\lambda}(a)$ ”, to give only a mild example of what the reader would otherwise be in for.)

First suppose that $\kappa \leq \lambda < \mu$ and $x \in \mathbf{M}_\lambda$. (In this paragraph, we suppress mention of the transition functions $f_{\kappa\lambda}$, for purposes of readability, preferring instead to write the incorrect “ $\lambda \models x \in a$ ” instead of “ $\lambda \models x \in f_{\kappa\lambda}(a)$ ”, to give only a mild example of what the reader would otherwise be in for.) If $x \notin \mathbf{M}^\mu$ then x gets into neither a nor $\text{Sep}_{a, \phi^{*f}, \mu}$, and we’re done. Otherwise, by the choice of μ ,

$$\lambda \models “x \in a \wedge \phi(x)” \text{ iff } \mathbf{V}_\mu[G] \models “\lambda \models “x \in a \wedge \phi(x)””.$$

The difference between \mathbf{M} (in $\mathbf{V}[G]$) and $\mathbf{V}_\mu[G]$ ’s version thereof is the partial order (ORD vs. μ) and the ordinal length of the inductive generation of terms (again, ORD vs. μ). The former is captured by ψ , the latter by f ; notationally, \mathbf{M} as interpreted in $\mathbf{V}_\mu[G]$ is $\mathbf{M}_{pre-\psi}^\mu$. So

$$\mathbf{V}_\mu[G] \models “\lambda \models “x \in a \wedge \phi(x)”” \text{ iff } \lambda \models_{\mathbf{M}_{pre-\psi}} “x \in a \wedge \phi^f(x)”.$$

By the previous lemma, the latter is equivalent to $\lambda \models "x \in a \wedge \phi^{f*}(x)"$. By the interpretation of Sep-terms, the latter is equivalent to $\lambda \models "x \in \text{Sep}_{a, \phi^{f*}, \mu}"$.

Now consider the case $\lambda \geq \mu$, $x \in \mathbf{M}_\lambda$. Since $\lambda \geq \mu > \text{dom}(a)$, if $\lambda \models "x \in a"$ then there are $\nu < \mu$ and $y \in \mathbf{M}_\nu$ such that $x = f_{\nu\mu}(y)$ and $\nu \models "y \in a"$. Without loss of generality ν can be chosen to be larger than the domains of a and ϕ 's parameters. For any such ν and y , by the same isomorphism argument as in the last section, $\nu \models \phi(y)$ iff $\lambda \models \phi(x)$. Furthermore, since $a \in \mathbf{V}_\mu[\mathbf{G}]$ (by the choice of μ), y is not only in \mathbf{M}_ν but also \mathbf{M}_ν^μ . Combining all of this with the definition of $f_{\kappa\lambda}(\text{Sep}_{a, \phi^{f*}, \mu})$ (which was highlighted just before the statement of the current theorem), we get the following equivalences:

$$\begin{aligned} \lambda \models "x \in a \wedge \phi(x)" &\longleftrightarrow \\ \exists \nu < \mu, y \in \mathbf{M}_\nu^\mu (x = f_{\nu\lambda}(y) \wedge \nu \models "y \in a \wedge \phi(y)") &\longleftrightarrow \\ \text{(by the case above: } \kappa \leq \lambda < \mu, \text{ with } \nu \text{ here playing the role of } \lambda) & \\ \exists \nu < \mu, y \in \mathbf{M}_\nu^\mu (x = f_{\nu\lambda}(y) \wedge \nu \models "y \in \text{Sep}_{a, \phi^{f*}, \mu}") &\longleftrightarrow \\ \exists \nu < \mu, y \in \mathbf{M}_\nu (x = f_{\nu\lambda}(y) \wedge \nu \models "y \in \text{Sep}_{a, \phi^{f*}, \mu}") &\longleftrightarrow \\ \lambda \models "x \in \text{Sep}_{a, \phi^{f*}, \mu}" & \end{aligned}$$

Strong Collection: In the presence of full Separation (see last paragraph), Strong Collection follows from Collection. Suppose $\kappa \models \forall x \in a \exists y \phi(x, y)$. Let λ, μ and f be as in the last paragraph. Then $\kappa \models "f$ is a bounding set for ϕ on $a"$.

Exponentiation: In order to highlight exactly why the following argument works for functions (yielding Exponentiation) but not for arbitrary relations (hence not yielding Subset Collection), it will be presented in terms of a relation R until exactly that moment when we need to use that R is not only a relation but also a function.

So let a and b be terms for sets in \mathbf{M}_κ , and p a forcing condition, R a term, and $\lambda \geq \kappa$ an ordinal so that $p \Vdash "\lambda \models R$ is a relation from a to $b"$. Since the property of being a relation between two given sets is Δ_0 , the assertion $p \Vdash "\lambda \models R$ is a relation from a to $b"$ is Δ_1 in \mathbf{V} . We will use later the same result for the property of R being a function from a to b .

Let H be the Σ_1 -Skolem hull of $\{a, b, \kappa, \lambda, R\} \cup \text{TC}(a) \cup \text{TC}(b) \cup \kappa$, of the same size as that set (where "TC" stands for transitive closure). Let ρ be the transitive collapsing function. Then

$$\rho^* H \models "\rho(p) \Vdash "\rho(\lambda) \models \rho(R) \text{ is a relation from } \rho(a) \text{ to } \rho(b)"".$$

ρ fixes p , a , and b because L_ω , $\text{TC}(a)$, and $\text{TC}(b) \subseteq H$. Furthermore, $\rho(\lambda) \leq \lambda$, so for some $\mu \leq \lambda$

$$\rho^* H \models "p \Vdash "\mu \models \rho(R) \text{ is a relation from } a \text{ to } b"" ,$$

and, by the same reasoning, for all $q \geq p$,

$$H \models "q \Vdash "\lambda \models \langle x_a, y_b \rangle \in R"" \text{ iff } \rho^* H \models "q \Vdash "\mu \models \langle x_a, y_b \rangle \in \rho(R)"".$$

However, the assertion “ $q \Vdash \mu \models \langle x_a, y_b \rangle \in \rho(R)$ ” is Δ_1 over both \mathbf{V} and ρ ”H (as the collapse of a Σ_1 substructure of \mathbf{V}), hence is absolute (between ρ ”H and \mathbf{V}), as is “ $q \Vdash \lambda \models \langle x_a, y_b \rangle \in R$ ” between H and \mathbf{V} . Therefore

$$q \Vdash \lambda \models \langle x_a, y_b \rangle \in R \text{ iff } q \Vdash \mu \models \langle x_a, y_b \rangle \in \rho(R).$$

Since $\mu \leq \lambda$

$$\text{if } q \Vdash \mu \models \langle x_a, y_b \rangle \in \rho(R) \text{ then } q \Vdash \lambda \models \langle x_a, y_b \rangle \in \rho(R).$$

Combining the last two assertions,

$$\text{if } q \Vdash \lambda \models \langle x_a, y_b \rangle \in R \text{ then } q \Vdash \lambda \models \langle x_a, y_b \rangle \in \rho(R),$$

i.e. $p \Vdash \lambda \models R \subseteq \rho(R)$.

Under the assumption merely that R is a relation, no more can be said. If, however, $p \Vdash \lambda \models R$ is a (total) function (from a to b), then, by analogous arguments to those above,

$$\begin{aligned} H &\models \text{“} p \Vdash \lambda \models R \text{ is a function”}, \\ \rho(H) &\models \text{“} p \Vdash \mu \models \rho(R) \text{ is a function”}, \\ p &\Vdash \mu \models \rho(R) \text{ is a function, and} \\ p &\Vdash \lambda \models \rho(R) \text{ is a function.} \end{aligned}$$

Since a (total) function cannot be properly extended to a function, $p \Vdash \text{“} R = \rho(R) \text{”}$.

By the cardinality constraint on H, there is an ordinal α (independent of λ and R) such that ρ ”H $\subseteq \mathbf{V}_\alpha$. In particular, $\rho(R) \in \mathbf{V}_\alpha$. Let g be the function with domain $\alpha \setminus \kappa$ such that $g(\nu) = \{ h \in \mathbf{V}_\alpha \mid h \text{ is a term naming an element of } \mathbf{M}_\nu \}$. Then g is a bounding set for ${}^a b$. By Separation (Restricted Separation suffices here), ${}^a b$ is a set.

– Subset Collection: Suppose that $\kappa \models \text{“for all total relations } R \text{ on } \mathbf{N} \text{ there is a total subrelation } R', R' \in F \text{”}$. There is an ordinal λ such that, for all terms G_μ in the transitive closure of F , $\mu < \lambda$; without loss of generality, $\lambda > \kappa$. Since for every such μ $\lambda \models G_\mu = \mathbf{N} \times \mathbf{N}$, $f_{\kappa\lambda}(F)$ can be evaluated in \mathbf{V} . By assumption, $\lambda \models \text{“there is a total subrelation of } G_\lambda \text{ in } f_{\kappa\lambda}(F) \text{”}$. This is, however, impossible, by the genericity of G , as follows. Let $\lambda \models \text{“} R' \in f_{\kappa\lambda}(F) \text{”}$. In \mathbf{V} , $R'(\lambda)$ can be identified with a total relation on \mathbf{N} (and so we will use the same notation $R'(\lambda)$ for this relation in \mathbf{V}). Let p be a forcing condition. Since p is finite and $R'(\lambda)$ total, let $\langle m, n \rangle \in R'(\lambda)$ be such that $\langle m, n \rangle \notin \text{dom}(p)$, and let q be $p \cup \langle \langle m, n \rangle, 0 \rangle$. $q \Vdash \text{“} R'(\lambda) \text{ is not a subrelation of } G \text{”}$, so $q \Vdash \lambda \not\models \text{“} R' \text{ is a subrelation of } G_\lambda \text{”}$. Hence the set of conditions forcing $\lambda \not\models \text{“} R' \text{ is a subrelation of } G_\lambda \text{”}$ is dense, and so in $\mathbf{V}[G]$ $\lambda \not\models \text{“} R' \text{ is a subrelation of } G_\lambda \text{”}$. This however contradicts the construction of λ , so the assumption on κ is false, and Subset Collection fails for relations on \mathbf{N} . ■

4 Reflection

Although this article is about Power Set and its variants, a result about Collection and its variants is right there, and so should be pointed out. Classically, over the other axioms of ZF, Reflection, Collection, and Replacement are equivalent. The proofs that Reflection implies Collection, which in turn implies Replacement are soft and easy proofs, and so carry over verbatim to intuitionistic logic. The reverse implications are less obvious and make essential use of classical logic; the standard proofs do not carry over. In fact, it was shown in [4] that Replacement does not imply Collection over the remaining axioms of IZF. It is natural to ask whether the same can be said of Collection and Reflection. Although this is still unknown, the first construction of this paper provides a model of CZF + Separation + \neg Reflection. The only axiom that's missing from IZF is Power Set, and Subset Collection is at least a partial substitute.

Moreover, the failure of Reflection is very strong. All true Π_1 facts are reflected in any transitive set, and any true Σ_1 sentence is reflected in any transitive set that contains a witness. In contrast, consider $\neg\forall x \forall y (x = y \vee x \neq y)$. This is true in the whole model, and false in every set.

References

- [1] P. Aczel, The type theoretic interpretation of constructive set theory, in A. MacIntyre, L. Pacholski, and J. Paris (eds.), *Logic Colloquium '77* (North-Holland, Amsterdam, 1978), p. 55-66
- [2] P. Aczel, The type theoretic interpretation of constructive set theory: choice principles, in A.S. Troelstra and D. van Dalen (eds.), *The L.E.J. Brouwer Centenary Symposium* (North-Holland, Amsterdam, 1982), p. 1-40
- [3] P. Aczel, The type theoretic interpretation of constructive set theory: inductive definitions, in R.B. Marcus et al. (eds.), *Logic, Methodology and Philosophy of Science VII* (North-Holland, Amsterdam, 1986), p. 17-49
- [4] H. Friedman and A. Scedrov, The lack of definable witnesses and provably recursive functions in intuitionistic set theories, *Advances in Math*, 1985, p. 1-13.
- [5] P. Martin-Löf, An intuitionistic theory of types: predicative part, in H.E. Rose and J. Sheperdson (eds.), *Logic Colloquium '73* (North-Holland, Amsterdam, 1975), p. 73-118
- [6] P. Martin-Löf, *Intuitionistic Type Theory* (Bibliopolis, Naples, 1984)
- [7] J. Myhill, Constructive set theory, *Journal of Symbolic Logic*, 1975, p. 347-382

- [8] M. Rathjen, The strength of some Martin–Löf type theories, *Archive for Mathematical Logic*, 1994, p. 347-385ff
- [9] M. Rathjen, The higher infinite in proof theory, in J.A. Makowsky and E.V. Ravve (eds.), *Logic Colloquium '95*, Springer Lecture Notes in Logic, Vol. 11 (Springer, New York, Berlin, 1998), p. 275-304ff